Sequential Bayesian Estimation of Parameters of Fission Processes

CASIS Workshop

May 13, 2015

Hema Chandrasekaran, James V. Candy, David H. Chambers, and Sean. E. Walston

Lawrence Livermore National Laboratory

LLNL-PRES-670603

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Background (1/2)

Goal

Rapid and accurate characterization of illicit special nuclear material (SNM) using neutron inter-arrival times measured by a passive neutron multiplicity counter.

Measurement

Time of arrival of neutron counts (from which inter-arrival times are calculated)

Neutrons from the correlated source (**blue**) arrive in short bursts. Neutrons from the random Poisson source (**red**) arrive in a steady stream.

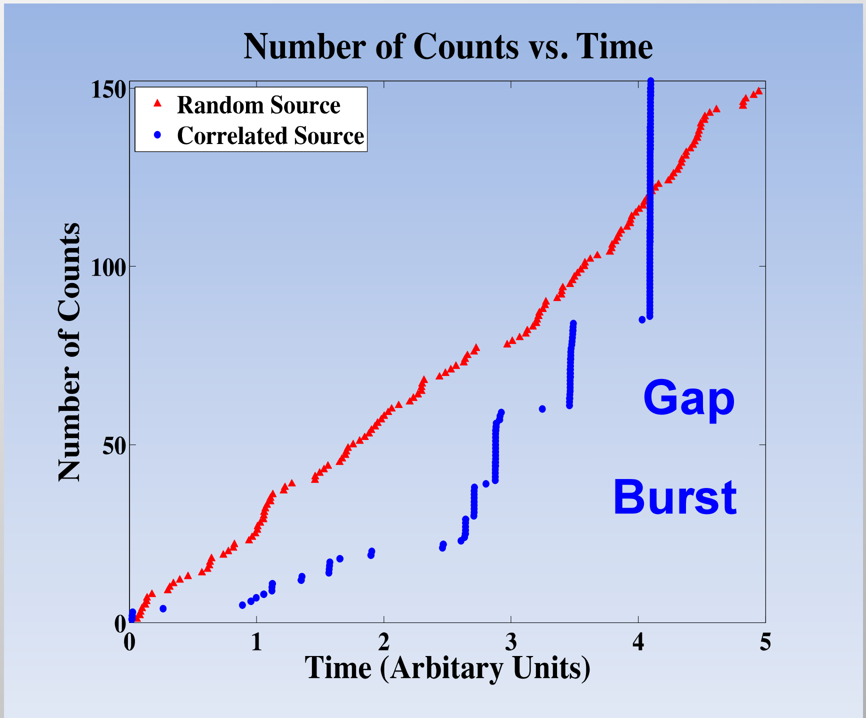

Background(2/2)

Defining characteristics of SNM

- Presence of neutrons – each fission creates on average 2-3 neutrons and neutrons arrive in short bursts followed by gaps
- Ability to sustain a fission chain, a series of fissions where each successive fission is induced by a neutron from a previous fission chain

Time Interval PDF as a Unique Signature of SNM

Theory provides an exact formula for the PDF of neutron inter-arrival times parametrized by fissile source mass, detection efficiency, system multiplication factor, neutron lifetime against detection.

This PDF, which is conditioned on a set of fissile source parameters, completely and uniquely characterizes the fissile source (quantity, and arrangement of the SNM)

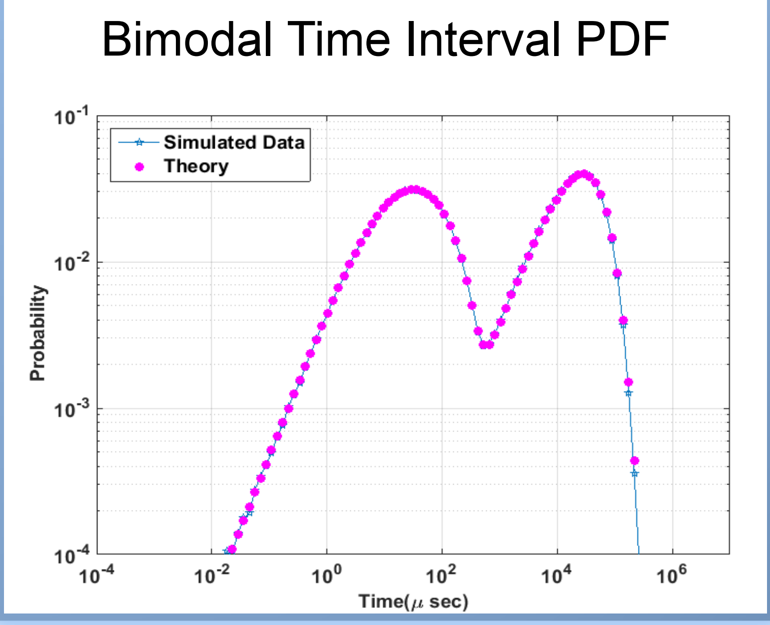

Parameter Estimation Problem

Problem Statement:

Given neutron inter-arrival times and the analytic PDF characterized by parameters of interest, estimate the parameters of the non-Gaussian, multimodal PDF as neutrons arrive.

Static and Independent Model Parameters

The parameters (source mass, detection efficiency, system multiplication factor, neutron lifetime against detection) are constants, uncorrelated and there is no drift or abrupt changes in model parameters over time.

Online or sequential estimation

Compute estimates of model parameters on-the-fly by continuously updating the posterior pdfs (hence Bayesian) of parameters as more observations become available in real time.

online algorithms are often more efficient — converge faster towards the target parameter values and need fewer computer resources.

Physics Model of a Fissile Source 1/2

$$\Pr[\tau|\Theta] = \underbrace{R_1 r_o n_o}_{Time\ between\ chain\ initiations} +$$

$$\underbrace{\frac{F_S}{R_1}\sum_{n=2}^{\infty} e_n(\epsilon)\left(\sum_{k=1}^{n-1} k e^{-k\lambda\tau}\right) b_o(\tau)\lambda}_{Time\ between\ neutrons\ in\ same\ chain}$$

R_1 is the count rate
F_S is the fission rate
r_o probability that no additional neutrons from the *same* fission chain are counted within a time τ
b_o is the probability that *no* neutrons are detected within the time interval τ
n_o is the probability of no additional neutron counts in time interval τ
τ inter-arrival time or time interval
$e_n(\epsilon)$ is the probability of detecting n neutrons from a fission chain

Physics Model of a Fissile Source 2/2

Embedded in the above equation is a set of various relations that capture the time-interval probability:

$$Fs = \frac{N_A}{A}\frac{\ln 2\ t_{1/2}}{t_{1/2}\ t_{1/2}^{SF}} m_S$$

$$R_1 = \epsilon\ q\ \mathcal{M}\ \overline{\nu}_S\ F_S$$

m_S mass of the source
λ is the inverse diffusion time scale
p is the probability that a neutron induces a fission
q is the escape probability (q = 1- p)
$\bar{\nu}_S$ is the average neutron count from a spontaneous fission
$\bar{\nu}_I$ is the average neutron count from an induced fission
A is the atomic number, N_A is the Avogadro Number

$\mathcal{M}$ is the system multiplicity and $\mathcal{M} = \frac{1}{1 - p\overline{\nu}_I} = \frac{1}{1 - k_{eff}}$

ϵ detection efficiency and $\epsilon = \frac{a}{b\mathcal{M}_e^2 + c\mathcal{M}_e}$ $\mathcal{M}_e = q \times \mathcal{M}$

$t_{1/2}^{SF}$ is the half-life

Sequential Bayesian Parameter Estimation (1/2)

Batch version: Estimate the posterior distribution $\Pr[\Theta\,|T_m]$ of source parameters Θ given the entire inter-arrival data set T_m up to time instant *m.* From Bayes' theorem we have

$$\Pr[\Theta|T_m] = \frac{\Pr[T_m|\Theta] \times \Pr[\Theta]}{\Pr[T_m]}$$

Sequential Version: The *posterior distribution* is given by

$$\Pr[\Theta|T_m] = \underbrace{\left(\frac{\Pr[\tau_m|T_{m-1},\Theta]}{\Pr[\tau_m|T_{m-1}]}\right)}_{W(\tau_m)} \times \underbrace{\left(\frac{\Pr[T_{m-1}|\Theta] \times \Pr[\Theta]}{\Pr[T_{m-1}]}\right)}_{\Pr[\Theta|T_{m-1}]}$$

Assume inter-arrivals are Markovian, that is, $(\tau_m, T_{m-1}) \rightarrow (\tau_m, \tau_{m-1})$.

Now the equation for sequentially propagating the posterior is

$$\Pr[\Theta|\tau_m] = W(\tau_m) \times \Pr\left[\Theta|\tau_{m-1}\right]$$ and $$W(\tau_m) = \frac{\Pr\left[\tau_m|\tau_{m-1},\Theta\right]}{\Pr\left[\tau_m|\tau_{m-1}\right]}$$

Sequential Bayesian Parameter Estimation (2/2)

Parameter vector Θ is a random constant with no associated dynamics …..

How then do you transition from $\Theta(\tau_{m-1}) \longrightarrow \Theta(\tau_m)$?

Set a non-informative prior on Θ using the known bounds on the parameters

$$\Theta(\tau_0) = \mathcal{U}[a,b] \quad \text{and} \quad R_{\Theta\Theta}(\tau_0) = \mathrm{Var}[\mathcal{U}[a,b]]$$

Constrained Random Walk

Try $\quad \Theta(\tau_m) = \Theta(\tau_{m-1}) + n_\Theta(\tau_{m-1})$

Until $\quad a \leq \Theta(\tau_m) \leq b$

where $\quad n_\Theta(\tau_{m-1}) \sim \mathcal{N}(0, \alpha^{m-1} R_{\Theta\Theta}(\tau_0))$

α is the annealing parameter and typically $0.9 \leq \alpha \leq 0.99$

Sequential Bayesian Parameter Estimation aka Particle Filter

What is Importance Sampling?

- Represent the PDF as a set of random samples (n)
- Approach an exact PDF as $n \to \infty$
- Obtain mean, covariance of the state vector PDF from the samples.
- Get a functional estimate of the PDF

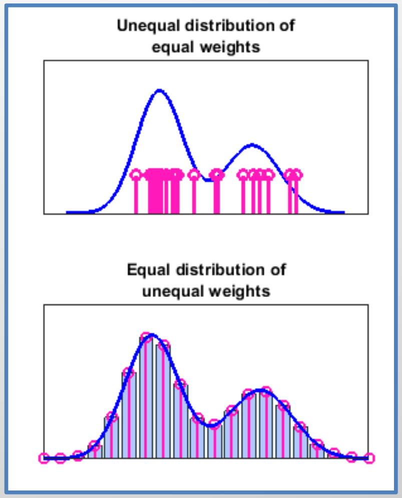

How do we go from prior to posterior?

- Update a sample from the prior to a sample from the posterior through the medium of likelihood function (Bayes' theorem in samples form)

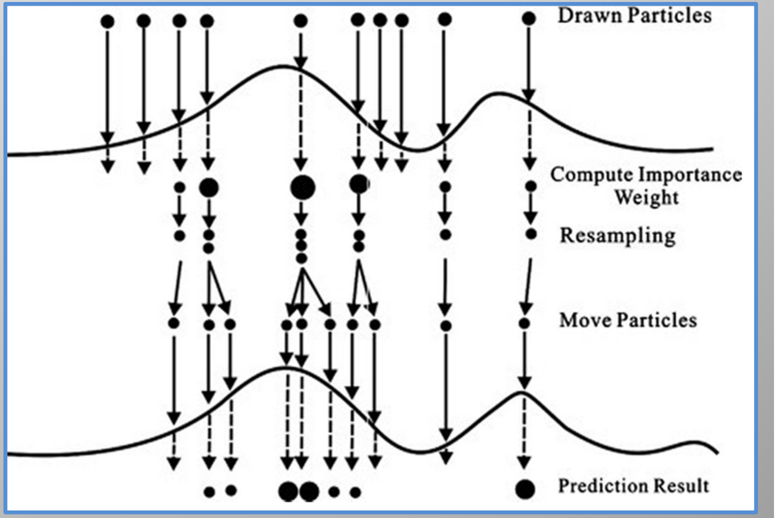

Why do we need to resample?

- Degeneracy phenomenon – particles collapse,
- all but one particle will have negligible weight
- Resample - eliminate particles that have small weights and generate a new set by resampling (with replacement) N times

Bootstrap Particle Filter
(with Constrained Random Walk and Residual Resampling)

Prior

$$\Theta_i(\tau_0) \sim \mathcal{U}[a, b]$$
$$W_i(\tau_0) = \frac{1}{N_p}, i = 1, 2, .., N_p$$
$$N_p = \text{Number of particles}$$

Prediction

Try $\Theta_i(\tau_m) = \Theta_i(\tau_{m-1}) + n_{\Theta_i}(\tau_{m-1})$

Until $a \leq \Theta(\tau_m) \leq b$

where $n_{\Theta_i}(\tau_{m-1}) \sim \mathcal{N}(0, \alpha^{m-1} R_{\Theta\Theta}(\tau_0))$

Update

$$W_i(\tau_m) = \Pr[\tau_m | \Theta_i(\tau_m)]$$

Normalize Weights

$$\mathcal{W}_i(\tau_m) := \frac{W_i(\tau_m)}{\sum_i W_i(\tau_m)}$$

Resample

$$\Theta(\tau_m) \Rightarrow \hat{\Theta}_i(\tau_m)$$

Posterior PDF

$$\hat{\Pr}[\Theta(\tau_m)|\tau_m] = \sum_i \mathcal{W}_i(\tau_m)\ \delta\Big(\Theta(\tau_m) - \hat{\Theta}_i(\tau_m)\Big)$$

Notional HEU Example (1/2)

Converges to the correct answers in ≈ 300 neutrons

- True mass = 25 kg
- Converged mass = 30 ± 9 kg
- True k_{eff} = 0.9
- Converged k_{eff} = .891 ± 0.026
- True diffusion time = 100 μs
- Measured = 91 ± 26 μs

3% Detection efficiency

No need for a priori knowledge of detection efficiency

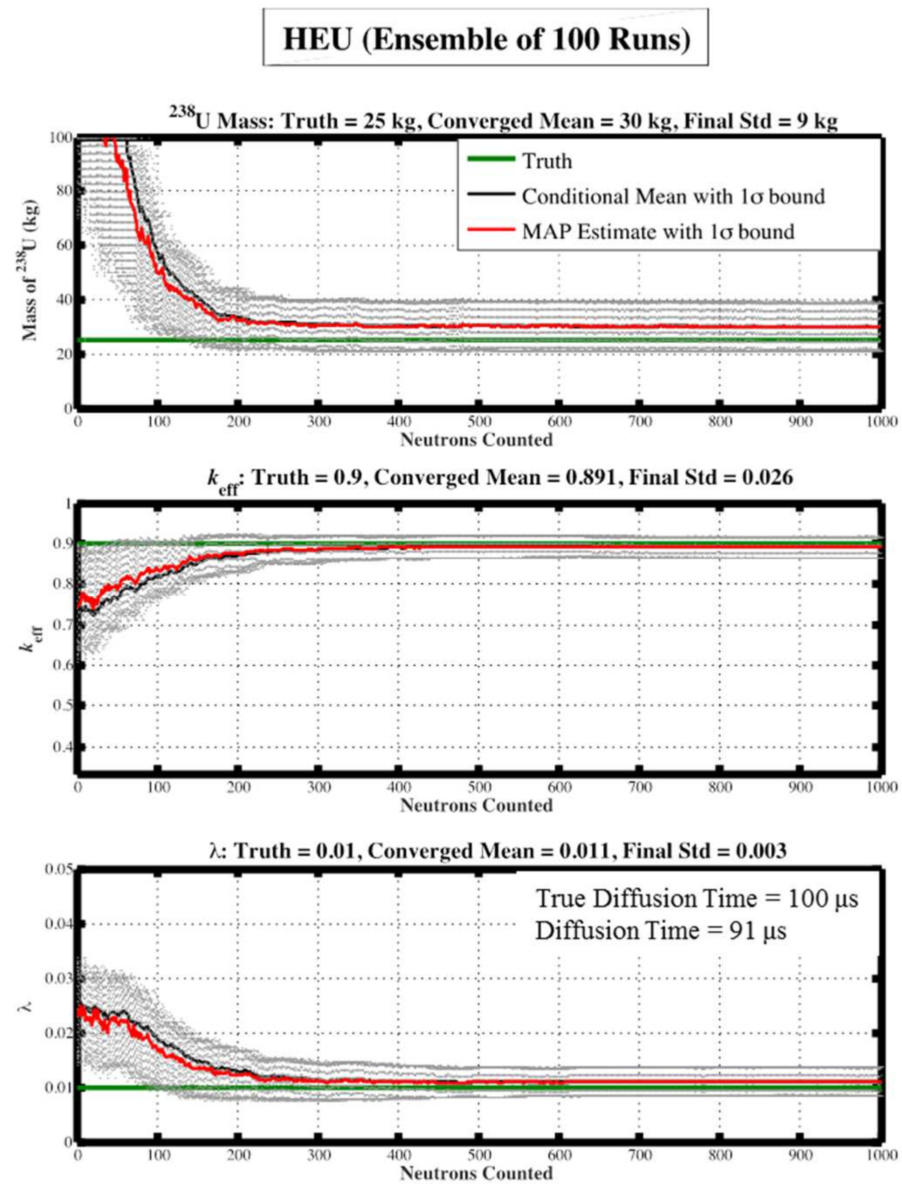

Notional HEU Example (1/2) – Convergence Results
(click on the figure to start movie)

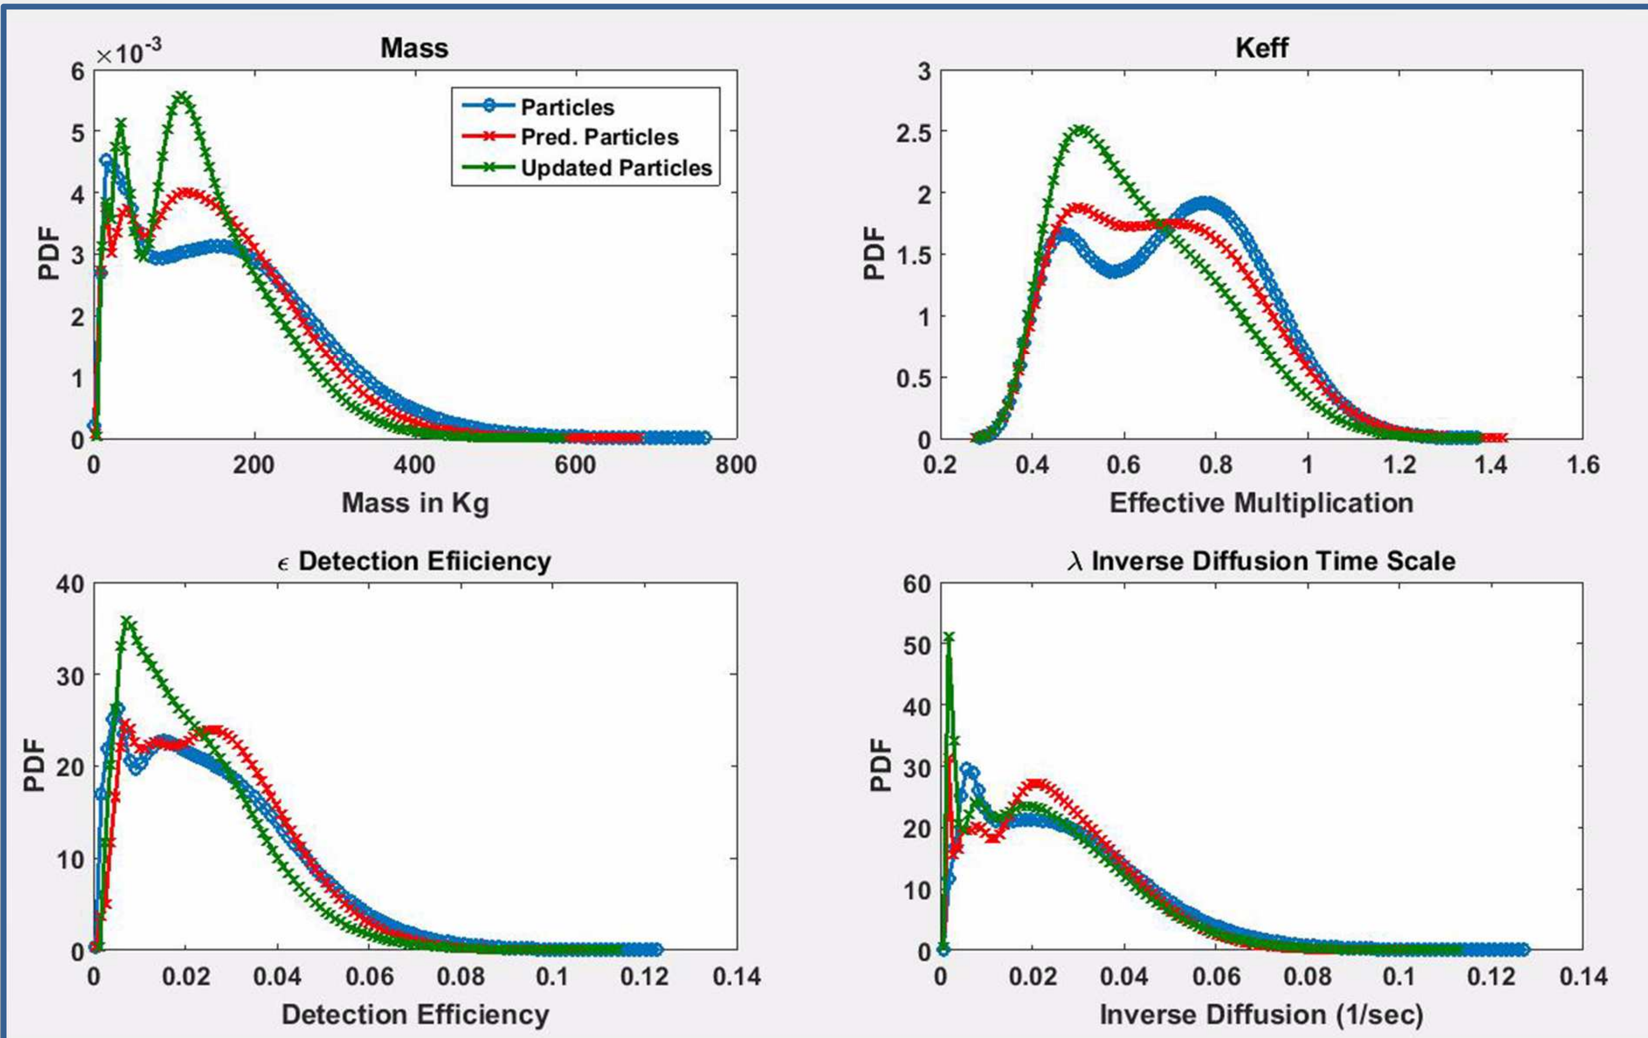

HEU in Steel (MCNP Model)

High Multiplication

Converges to the correct answers in ≈ 400 neutrons

- True mass = 1.837 kg
- Converged mass = 3 ± 6 kg
- True k_{eff} = 0.9512
- Converged k_{eff} = .932 ± 0.016
- Diffusion time = 37 ± 10 μs

1.5% Detection efficiency

No need for a priori knowledge of detection efficiency

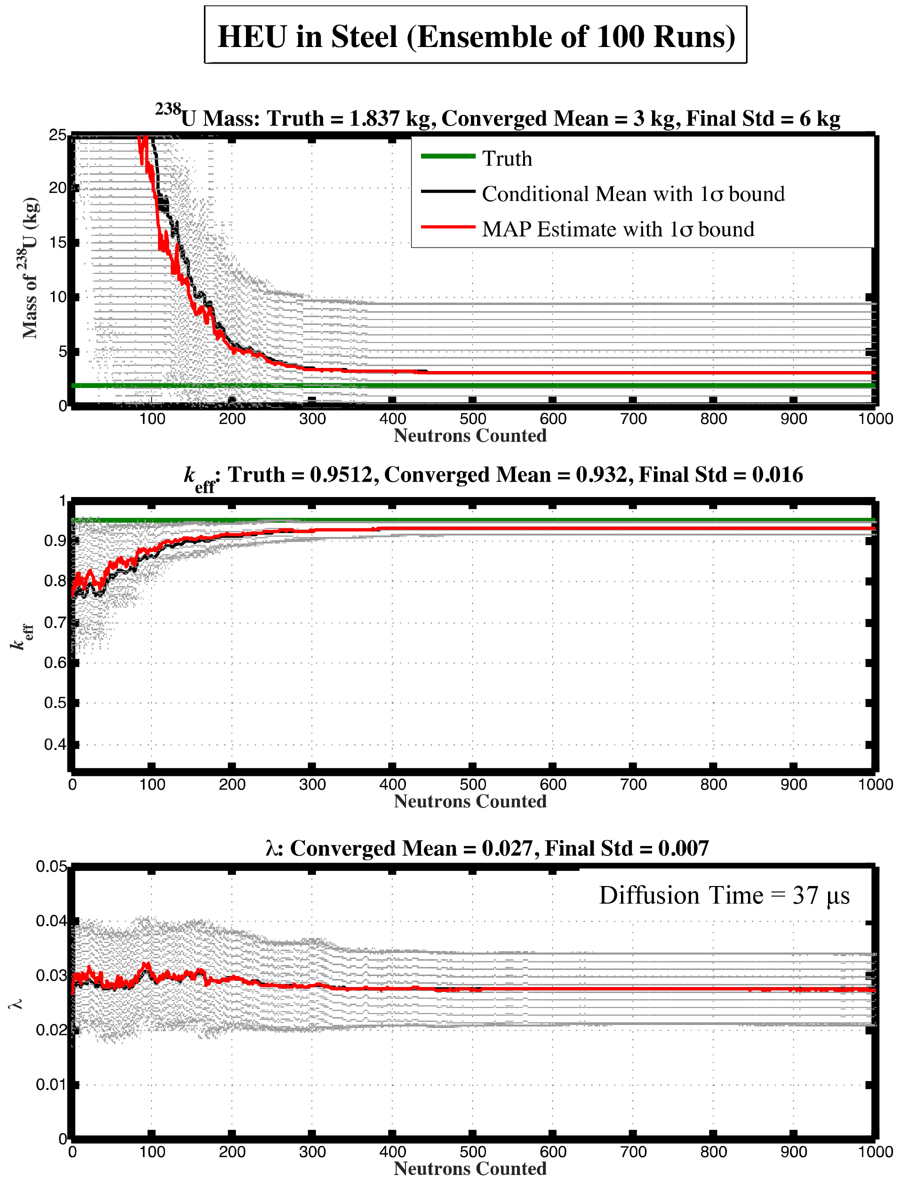

HEU Shell in Poly (MCNP Model)

Low Multiplication

Converges to the correct answers in ≈400 neutrons

- True mass = 0.3962 kg
- Converged mass = 1.5 ± 0.5 kg
- True k_{eff} = 0.5331
- Converged k_{eff} = 0.59 ± 0.08
- Diffusion time = 71 ± 29 μs

2.5% Detection efficiency

No need for a priori knowledge of detection efficiency

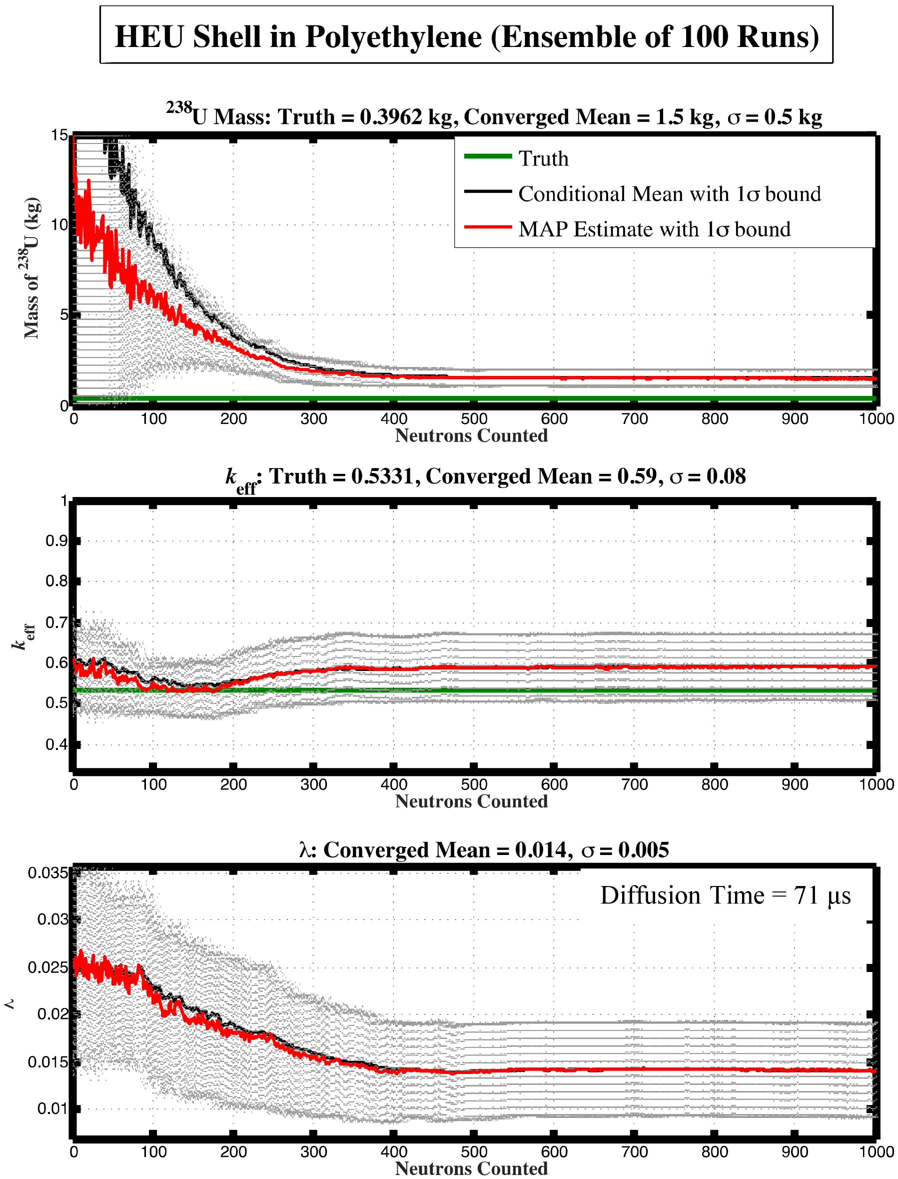

Convergence, Accuracy, and Local Maxima

Theory guarantees convergence of the empirical distributions generated by particle filters toward the true distributions as the number of particles (N)$\rightarrow\infty$.

Computational burden increases as N increases. Using smaller N leads to bias in the estimate of the posterior mean as PF can get trapped in local maxima. One way to combat bias would be to iteratively narrow the 4D parameter search space while increasing N, the number of particles used.

To obtain uncertainties on the estimated parameters, start simultaneously multiple PF runs with random starting points and take the standard deviation of the converged results.

Summary

The problem of estimating the parameters of a fission process is stochastic since each neutron arrival, as it is transported through its path to the detector, is a random draw from the underlying time-interval distribution.

It requires a sequential Bayesian approach to provide MAP estimates of the parameters at each time step.

We applied an analytical likelihood PDF (Snyderman and Prasad, 2012) to construct an MCMC particle filter capable of estimating the source parameters and their accompanying probability distributions.

We demonstrated that we could reliably estimate the parameters of a fission process using MCNP simulation data.

References

1. J. V. Candy, S. E. Walston, and D. H. Chambers, “Sequential Detection of Fission Processes: A Physics-Based Approach,” submitted to IEEE Transactions on Nuclear Science.
2. D. H. Chambers, H. Chandrasekaran, and S. E. Walston, “Fourier method for calculating Fission chain neutron multiplicity distributions,” Lawrence Livermore National Laboratory Report, LLNL-JRNL-668317-DRAFT
3. M. Prasad and Snyderman, N., “Statistical theory of fission chains and generalized Poisson neutron counting distributions,” Nuclear Science and Engineering, Volume 172,,Number 3, November 2012, Pages 300-326.
4. M. Prasad, Snyderman, N., Verbeke, J., and Wurtz, R. “Time interval distributions and the Rossi correlation function,” Nuclear Science and Engineering, Volume 174, Number 1, May 2013 , Pages 1-29.
5. S. Walston, “A guide to the statistical theory of fission chains,” 2Ed., Lawrence Livermore National Laboratory Report, LLNL-TR-584832,2012.
6. J. V. Candy, Bayesian Signal Processing, Classical, Modern and Particle Filtering, Hoboken New Jersey: John Wiley/IEEE Press, 2009.
7. N. Gordon, D. Salmond, and A. F. Smith, “Novel approach to nonlinear/non-Gaussian Bayesian state estimation,” IEE Proc. F, Radar Signal Process., vol. 140, pp. 107–113, 1993.
8. A. Doucet, N. De Freitas, and N. Gordon, Eds., Sequential Monte Carlo Methods in Practice. New York: Springer, 2001.
9. O. Cappé and E. Moulines, “On the use of particle filtering for maximum likelihood parameter estimation,” in European Signal Processing Conference (EUSIPCO), Antalya, Turkey, September 2005.
10. Lang, L., Chen, W., Bakshi, B. R., Goel, P. K., and Ungarala, S. (2007). Bayesian estimation via sequential Monte Carlo sampling—Constrained dynamic systems. Automatica, 43(9), 1615-1622.
11. ReBEL-0.2.6, a Matlab toolkit for sequential Bayesian inference in general state space models, by Rudolph van der Merwe and Eric A. Wan.

This work was supported by the U.S. Defense Threat Reduction Agency (DTRA), under DTRA10027-10273

Backup Slides

Physics Model of a Fissile Source 3/3

The probability of creating n neutrons of the ν emitted with probability P_ν is given by

$$e_n(\epsilon) = \sum_{\nu=n}^{\infty} P_\nu \binom{\nu}{n} \epsilon^n (1-\epsilon)^{\nu-n}$$

The following probabilities complete the distribution

$$r_o = \frac{F_S}{R_1} \sum_{n=1}^{\infty} e_n(\epsilon) \Big(\sum_{k=0}^{n-1} e^{-k\lambda\tau} \Big)$$

$$b_o(\tau) = \exp\Big[- F_S \int_0^{\tau} \frac{1}{1-e^{-\lambda t}} \Big(1 - \sum_{\nu=0}^{\infty} P_\nu (1-\epsilon(1-e^{-\lambda t}))^{\nu} \Big) dt \Big]$$

$$n_o(\tau) = r_o(\tau) \times b_o(\tau)$$

Lawrence Livermore
National Laboratory